ALERT LOGIC®

CLOUD SECURITY REPORT

Research on the Evolving State of Cloud Security

2017


ALERT LOGIC®

CONTENTS

INTRODUCTION

OUR REPORT BY THE NUMBERS

550 DAYS

(AUGUST 1, 2015 – JANUARY 31, 2017)

2,207,795

TOTAL TRUE POSITIVE SECURITY INCIDENTS ESCALATED BY ALERT LOGIC ANALYSIS

32.5 MILLION

EVENTS ASSOCIATED WITH INCIDENTS

147 PETABYTES

OF DATA ANALYZED

3807 CUSTOMERS

ANALYZED

In the last few years the IT industry has crossed the chasm and cloud adoption no longer looks like an exotic proposition. This is as profound a paradigm shift as the Internet transformation appeared to be two decades ago, and it is driving an equally powerful change in the way we must evaluate the threat landscape. In 2017, we see a consolidation of threats in the very topmost layers of the computing model. The shift suggests that new approaches and fresh thinking will be required for businesses looking to increase their security posture and manage risk in cloud and hybrid environments.

Our data indicates that currently, web applications are the soft underbelly of your organization – if only by process of elimination.

After years of refinement, cloud service providers (CSPs) are expert at securing the lower (physical, logical, network) layers of the stack. Even attacks a bit higher, at the OS level, are on the decline – or stymied by the speed at which CSPs can apply patches and updates. That leaves the upper reaches of the process; the application components of your stack.

In a world where robust programs can burst into life in days or weeks, applications are asked to do a lot – handle inputs from millions of users simultaneously, interact with data stores around the globe, process and return results in milliseconds, and look good (and personalized) doing it. In that environment, developers don't tend to hear the nervous clucking of security professionals as much of a siren song.

Adding to the excitement, as web-based applications have gained in popularity, they have moved into (or are native to) cloud environments. This puts interesting new pressures on organizations, which are ultimately responsible for keeping their data secure – but increasingly unable to exercise fine control over the apps they run, their patch-application schedule, and so forth.

Finally, even though major attacks and zero-day vulnerabilities continue to make the headlines, attackers are increasingly happy to fly below that kind of radar. Their movement toward assembling a chain of vulnerabilities to build stealthy, resilient attacks is accelerating. Our data shows that many breaches today happen via lateral movement,

in which attackers exploit vulnerabilities in less-critical assets then make their way to the true target. This means that though individual security incidents are still interesting in and of themselves, it's often only when they're grouped and analyzed together that the whole turns out to be far greater than the sum of its parts.

Smart attackers, always seeking the weakest spots in network defenses, understand the changing attack scene and have retooled accordingly. This Cloud Security Report looks primarily at web application attacks, which account for 75 percent of all the incidents we flagged in our 18-month evaluation period. Such attacks affected 85 percent of all Alert Logic customers, with injection-style attacks such as SQLi (SQL injection) currently dominant. Our observations jibe with other findings throughout the industry. Industry observers such as Verizon[1], Gartner[2], SANS[3], and such state that vulnerable applications are the number-one means by which attackers breach data; meanwhile, according to our colleagues at Veracode, 56 percent of all PHP apps alone had at least one SQLi vulnerability[4].

We identify four other significant categories of attack methods targeting our clients: brute force attacks, malware infections, undesirable outside reconnaissance, and denial of service attacks. Of these four, we find high levels of activity in the brute force and server-side malware categories; this report also examines those more closely. Recon is also represented in our statistics as observable activity, though it's mainly of interest to our own work at identifying patterns of threat – we track and learn from it, but the excitement is elsewhere.

We found that defenders in very different sectors have much more in common than they might realize. Our research shows that though attack methods rise and fall in popularity, verticals saw generally similar attack methods over the time period covered in this report, raising interesting questions about attacker groups sharing tactics and even tools. We focused on five noteworthy verticals – Financial Services and Insurance; Health Services; Information Technology and Services; Production, Manufacturing, and Logistics; and Retail and Accommodation – to pinpoint attack vectors and patterns targeting those sectors.

KEY TAKEAWAYS

- **Web applications are the soft underbelly of your organization – the number-one means by which attackers breach data.**
- **The movement toward assembling a chain of vulnerabilities to build hard-to-detect, resilient attacks is accelerating.**
- **Hybrid networks, with portions scattered among public clouds, private clouds, and on-premises systems, are at greatest risk.**
- **Organizations in different sectors suffer from very similar attacks – and can learn much from each other.**

Finally, we spotlight two high-profile security incidents that occurred (or came to light) just after our early-2017 data window – one an example of a security incident handled very well, and the other an example of a security incident handled very poorly. While neither organization is an Alert Logic® customer, both offer useful practitioner lessons.

Who can benefit most from reading this report? Organizations working through decisions about rebalancing their systems among public, private, and on-premises solutions can see from our findings where to put extra effort into securing critical functions. Network and application architects who need to understand where danger lies can learn from the experience of those who have triumphed, or not, over the changing nature of enterprise threats. And those charged with inspiring their organizations to do better and be better at understanding current threats will find plenty of guidance as to how to proceed for maximum effect.

EXECUTIVE INSIGHT

Why is the sky blue? What is the meaning of life? Why did the chicken cross the road? Is the public cloud really less secure than on-premises data centers? No one has answers to many of these eternal questions, but we can shed some light on that last one – we have no indication that public cloud is less secure. In fact, there is an increasing body of evidence to the contrary.

For several years, we have observed that across the industry, security incident rates in public cloud environments are lower than they are on-premises. Though we have chosen not to highlight this in past Cloud Security Reports, we've confirmed this perception over time by close analysis of our own data. With years of observations and a clearly established pattern in hand, we are now confident in concluding that public cloud environments have lower observed incident rates than on-premises data centers.

To be fair, our data set does not conclusively prove that public clouds are "more secure," if such a question can even be framed effectively or answered definitively. But we do know that within our customer base, we less frequently see malicious activity in public cloud environments, even though web applications are one of the most dominant workloads there. And we know that web apps account for the highest share of attacks leading to breaches. Aside from incident rates, the most conclusive way to answer the question of which infrastructure model is most secure would be to focus strictly on breach data – a topic for future research.

We have anecdotal evidence that there are good reasons for this disparity in incident rates, as well as several factors which may skew our results to some degree. On average, the on-premises data centers Alert Logic monitors are larger than typical public cloud environments, with a larger number of assets. They are often directly attached campus networks, which aggregate many user endpoints on network links we monitor in addition to servers. It thus stands to reason we'd see higher incident rates on-premise.

Meanwhile, though on-premises environments see more incidents per customer than do those in the public cloud, we have noted even higher incident rates in hybrid cloud deployments. As interesting as this particular data point might be at a glance, we urge caution – while we're very confident in our public vs. on-premises classification, extending the analysis to hybrid is complicated by the fact that the industry cannot agree on what we all mean by "hybrid IT" or "hybrid cloud." However, it's possible that installations combining public and on-premises components catch the worst of both worlds – not as lockstep in receiving updates as all-public installations, not as carefully attended as on-premises installations with dedicated staff. We'll be watching this potential trend closely for our next report.

Some of the reviewable disparity in incident rates between different types of installations is, however, directly ascribable to the differences between public cloud and on-premises solutions. Those differences, we believe, account for lower incident rates in public cloud installations. Two differences are especially important:

- The significant pattern we see emerging time and again in public cloud installations is application-level segmentation of infrastructure. The best cloud administrators we know tend to segment each application in its own VPC (Virtual Private Cloud), which dramatically lowers the blast radius of any single breach. Even the smallest WordPress or Drupal apps get their own VPC, so there is less opportunity for attackers to move laterally, or to launch attacks able to unfold rapidly into enterprise-wide calamities.

- Application-level segmentation does not solely account for the decrease in incident rates in our customers' public cloud environments, but the level of control public clouds offer, in the form of better security mechanisms and easier administration, tends to translate into a smaller attack surface. We know this from our own experience – using public cloud resources allows us to manage our infrastructure better, thus more securely.

"We are now confident in concluding that public cloud environments have lower observed incident rates than on-premises data centers."

51%

HIGHER RATE OF SECURITY INCIDENTS IN ON-PREMISES DATA CENTERS

What does it all mean? Even with years of data in hand, it's too early to draw major conclusions. All we know is that we tend to see fewer incidents in public clouds than anywhere else. In future CSRs, we hope to look at this trend from other angles, which may reveal more insights.

For the moment, even the public cloud is not so secure that there is nothing to worry about; far from it. While we saw a 51% higher rate of security incidents in our customers' on-premises data centers, this still leaves each public cloud deployment to withstand just over (on average) around 400 incidents in the 18-month period covered by this report. And even lower incident rates do not necessarily translate to lower risk – especially when, as is increasingly more common, businesses rely on the public cloud to handle their highest-value assets. (Interestingly, we see little difference in incident rates among the public cloud providers; rates across AWS and other cloud providers are within the 2%-3% margin of error.) In any case, "what, me worry?" is not a feasible security stance; any enterprise that treats public clouds' lower incident rate as an invitation to be less diligent about security is seriously asking for trouble. It makes no sense to move to a safer neighborhood and then proceed to leave the doors of your house wide open!

We know from years of observation and experience that it's not the number of attacks or vulnerabilities that matter, but the attacks and vulnerabilities that ultimately lead to a breach. And when it comes to the biggest drivers for breaches, web apps are the dominant driver across the board, surpassing even the likes of privilege misuse, point-of-sale compromises, and exploitation of unpatched operating system vulnerabilities. Wherever your data is hosted, this core truth should drive your response and your prioritization of security issues.

We're not huge fans of sweeping conclusions, but if you must, here's ours: Go ahead and deploy more assets in public clouds, but continue to follow best security practices when you do. In fact, you should feel comfortable accelerating these migrations. The security benefits of public clouds are now both qualitative and quantifiable.

Misha Govshteyn

Senior Vice President, Products & Marketing and
Co-founder of Alert Logic

METHODOLOGY

HOW WE BUILT THIS REPORT

KNOW YOUR MAYHEM

EVENT

Alert Logic defines an event as a finding fired by a security control as a result of detected suspicious or malicious behavior.

INCIDENT

is an event or group of events that have been confirmed as a valid threat warranting further investigation, analysis, and possibly response.

BREACH

is an incident that results in the confirmed loss of data to an unauthorized party.

EXPLOIT

is a tool or technique, often embodied in code, that takes advantage of (exploits) a vulnerability such as an unpatched software bug or a misconfiguration.

ATTACK

is the use of exploits or other means to harm, steal from, block access to, or otherwise wrongly impede a computing resource.

Our dataset is sourced from 32.5 million events and over 2.2 million verified security incidents captured in Alert Logic network intrusion detection systems between August 1, 2015 and January 31, 2017. We employ a patented incident analytics platform that evaluates multiple indicators of compromise in event data to identify and classify incidents. Our security analysts further assess indicators of compromise within the context of relevant threat intelligence, coupled with customer asset blueprints, to validate incidents. In some situations, machine learning techniques are also deployed, either to refine analysts' conclusions or to spot certain kinds of multi-stage attacks. The result is low-noise, verified, actionable incident data that provides the context customers need to take action.

The customers in our report dataset represent a broad range of industries (452 unique SIC codes) and organization sizes (from 100 to 10,000+ employees). The majority of customers are securing cloud workloads – 82% have workloads hosted on an IaaS (Infrastructure as a Service) and/or Hosted Private Cloud–and that ratio has remained relatively steady over the last three reporting periods.

In addition to their public cloud footprints, about a third of our customers maintain on-premises or hybrid infrastructure. This percentage is a bit lower than that of the world as a whole, in which fully three-quarters of networks are held to be some form of cloud-prem hybrid. While as a rule we don't have visibility into the details of all applications and workloads our customers are running, on-premises infrastructures typically support legacy or home-grown business applications and workloads that aren't easily portable to public clouds.

In previous reports, a large portion of our analysis was organized by industry. While industry can be an indicator of security and compliance requirements, it does not provide a full picture. IT and security teams are not only charged with administering and securing the industry-aligned workloads that are core to their company's external business, but they are also looking after the workloads that support the back office and internal decision-making functions.

To provide you with a deeper and more relevant level of analysis and insight, we are trying something new for this edition of the report. Instead of segmenting our analysis by industry and primarily looking for trends and insights within an industry group, we're focusing our analysis against a prioritized set of incident types, and the workloads and environments most at risk.

Finally, for ease of reading we have been very selective in the charts we present in the body of this report. For a more detailed look at the data underpinning our conclusions, please see the tabular data in Appendix A.

THE FINDINGS

AN OVERVIEW OF ATTACK TYPES AND TARGETS

While our data is categorized by incident type, we know from years of analysis that individual security incidents are often key milestones in a more complicated attack progression. To that end, we remind readers not to think of attacks as simply an indicator that one kind of event has taken place, even when incidents are categorized for report purposes as being of this or that event type. Combinations of events yoked together to formulate incidents, which in turn formulate attacks, are the order of the day. That said, the easiest way to think about defending your network is to think about the incidents that touch it.

SECURITY INCIDENT TYPES OBSERVED AND ESCALATED

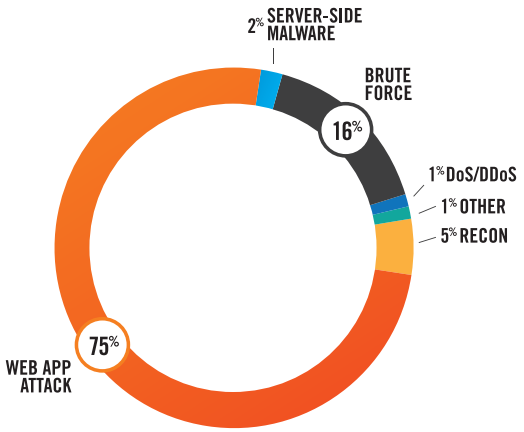

Figure 1: Our security incident data indicates that over three quarters of all the events we saw during the reporting period involved web application attacks – which includes both recon and exploits targeting web applications.

The chart above shows that attacks on web applications tower above all other threats to our customers' assets. Brute force attacks, which rely less on a fine understanding of the target, make up a nontrivial one-sixth of incidents in our sanitized dataset, and together the two account for nine-tenths of all attempts. We look at each type of attack in greater depth below.

Even though an active malware attack can run a network administrator ragged and get broad attention from the general public, we found that they were a relatively small percentage of incidents in our dataset. This does not indicate that malware isn't a problem, but that anti-malware protections are continuing to fight against a heavy onslaught, and that many of the incidents we saw that related to malware were actually tracking the results of an infection (most notably the calls back to a command-and-control server) rather than the end user end-point infection. However, change is in the air, with various new forms of malware that rely on little to no end user interaction gaining ground. We are observing an increasing number of server-focused malware payloads that use web applications and cloud storage services for infection, command-and-control functionality, and distribution. We discuss current findings and future predictions in greater depth below.

Recon incidents observed by Alert Logic include scans from legitimate security vendors and malicious bad actors. We've stripped the known-good vendor scans from the sanitized chart above as these are actually benevolent forms of reconnaissance; for instance, a vulnerability assessment test run against your systems will cause events that on first glance appear to be recon. Alert Logic is able to identify known-good traffic events such as these and manage its escalations accordingly. The other kind of recon accounts for about five percent of the incidents observed during the period of this report. Scans indicate that someone or some tool (e.g. Nessus or Metasploit) is eyeing your network – possibly seeking attack-ready weaknesses. We consider recon a datapoint that may indicate the very beginnings of an external attack. By itself, the data is not actionable, but systems such as Alert Logic's can evaluate and treat apparently isolated incidents such as recon as early warnings, giving defenders a crucial heads-up when it counts.

Finally, there's DoS and DDoS. Denial of service events account for just one percent of incidents observed during the period of this report. DoS and DDoS (distributed DoS) attacks differ mainly by the number of events observed and the number of IP addresses involved in the attack.

For all incident categories, we note that specific vulnerabilities observed and recorded well over a decade ago are still being exploited today. (Attackers like tried-and-true vulnerabilities and exploits that, thanks to sloppy patching habits, continue to work well.) Of the total incidents in our database, we found that over 70% of them were related to vulnerabilities reported in 2014 and 2015. However, we saw vulnerabilities representing the full 29-year history[5] of modern malware, including a lively 4% of incidents traceable to 1999-era Windows IIS vulnerabilities – yes, vulnerabilities old enough to vote.

The strong showing of 2014-era vulnerabilities is due to Shellshock, which was disclosed in September of that year. Shellshock targets Internet-facing services including certain web server deployments, and uses Bash to process certain requests, allowing an attacker to cause vulnerable versions of Bash to execute random commands. Heartbleed, discovered in 2012, with the fix introduced in 2014, is similarly well-represented in the dataset. It exploits a vulnerability in the OpenSSL cryptography protocol. (We should note that the very notoriety of certain major vulnerabilities, and the interest that security researchers take in their prevalence and spread, cause the vulnerabilities themselves to be even more prominent in scanning "noise" around the Internet and in our data.) And ancient vulnerabilities aren't just useful for attacking web applications; most attempted brute-force attacks used vulnerabilities dating from 2000 and 2001.

Now that we're clear on attack types, what can be said about the targets? Our analysis found that while types are consistent whether the targeted installation is hosted on a public cloud, in a traditional on-premises network, in a privately hosted cloudspace, or in some combination of those options, app attacks rule the day. We did, however, notice a curious correlation to hosting environment type and attack frequency.

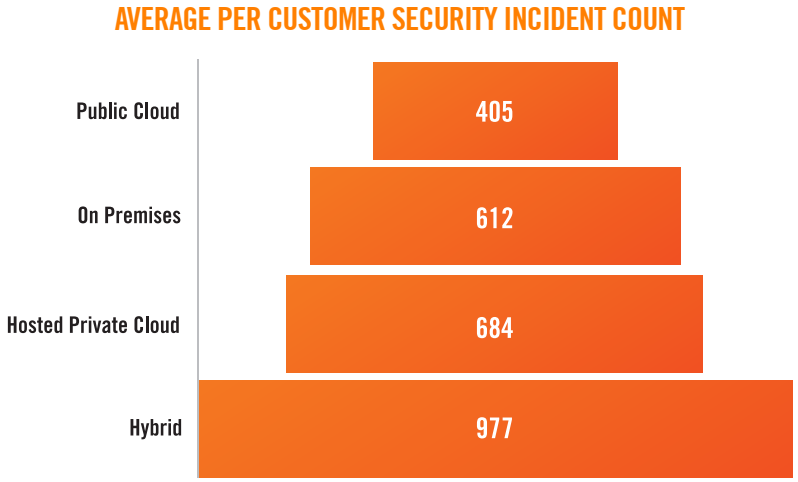

Figure 2: Our data over the 18 months of our report indicates that hybrid networks experience, on average, over twice as many security incidents as public cloud installations. Average per customer incidents are calculated based on the total incidents and total active customers observed during the 18 month analysis period. The customer population was not limited to only customers present for the entire 18 month period. Some variation could apply.

Our data shows a remarkable 141% higher rate of incidents per customer for hybrid installations – that is, those who have some combination of public cloud, on-premises network, and/or hosted private cloud in use – than we did for purely public cloud installations. On-premises installations experienced about 69% more security incidents per customer than did for enterprises relying strictly on public cloud services, while hosted private cloud entities saw about 51% more.

A couple of observations piqued our interest and are apt to be the subject of further investigation in our next report. Most crucially,

Figure 3a: These two charts, and the two below them, show an intriguing difference in the number of customers affected by most kinds of security incidents. We'll talk more about SQLi, the glaring exception, later in this report.

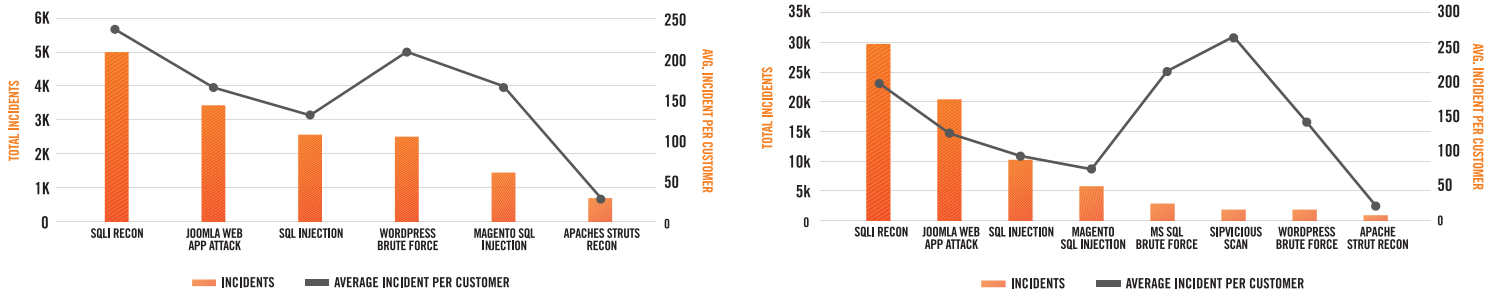

Figure 3b: These two charts, and the two above, show both similarities and differences in incident types and average number of incidents per customer among public cloud, on-premises, hybrid, and hosted private cloud environments. Note that all four environments share the top three most common incident types: SQLi reconnaissance-related activity, Joomla Web applications attacks, and SQL injection (SQLi) issues.

the industry definition of "hybrid clouds" remains fluid, which means that sweeping statements about the security of mix-and-match solutions aren't quite possible without years of further observation. Still, the comparison is at this point very difficult to ignore.

We noted also that attacks on a single entity across multiple clouds more closely resemble a targeted attack than an attack of opportunity – a tantalizing glimpse of an attack pattern we look forward to investigating further. Targeted attacks of this sort include malicious code designed to get and retain control of an environment. Our data showed examples of this attack pattern using command-and-control-reliant malware such as GhostRAT and Ranky.

KING OF THE HILL : WEB APPLICATION ATTACKS

Web application code is, quite frequently, the vector preferred by attackers to gain unauthorized access, compromise systems, and exfiltrate data. The avenue that makes the Internet so effective for the exchange of ideas, information, and e-commerce is also the very thing that makes it vulnerable, particularly when it comes to ports 80 (http) and 443 (https). Any device with the ability to communicate on the Internet over these ports can access web servers and gain access to information and data – sometimes far more information and data than the host intends.

TOP 6 WEB ATTACK TYPES OBSERVED

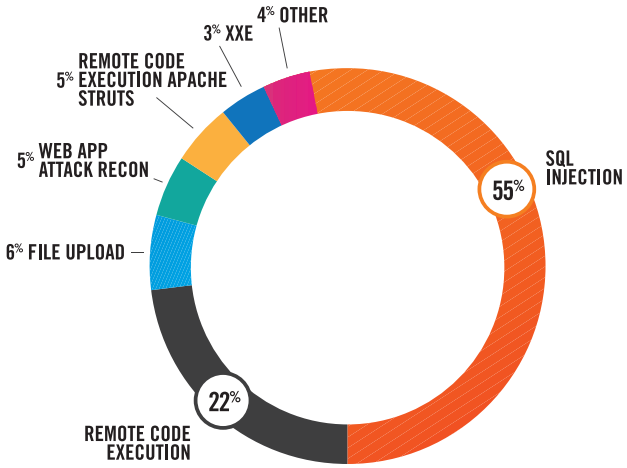

Figure 4. Four-fifths of OWASP Top 10 application attacks escalated were down to two well-understood methods: SQL injection and remote code execution. DoS/DDoS exploits are observed but the focus of this analysis is on exploits observed through network packet inspection. Brute force is discussed separately due to its significance.

In the Verizon 2016 Data Breach Investigation Report[6], researchers noted that "web applications account for only 8% of overall reported incidents (whether they were successful or not), but attacks on web applications accounted for over 40% of incidents resulting in a data breach, and were the single-biggest source of data loss." Our data confirms the dominance of application attacks in the wild; they account for 75% of verified security incidents, and affected 49% of our analyzed customers.

WEB ATTACK INCIDENTS PER MONTH

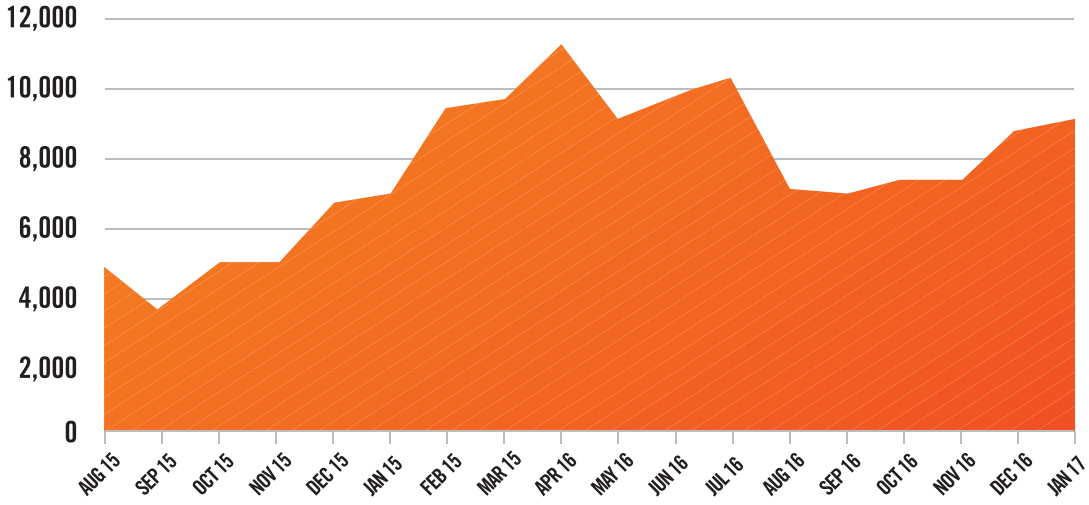

Figure 5: If it feels as if web application attacks have risen overall in the last year and a half, you're not wrong.

Our data further indicates that SQL injection was the attack vector used most frequently by far on customer environments, with 55% percent of the incidents we saw falling into that category. Injection-style attacks are fascinating – and extremely difficult to avoid, thanks to the very nature of modern data-driven applications. (They're not simply a SQL-related problem either, as our section below explains. In addition, e-commerce platforms such as Magento and Drupal are often targeted as part of SQLi campaigns, above and beyond standard SQL-driven web applications.)

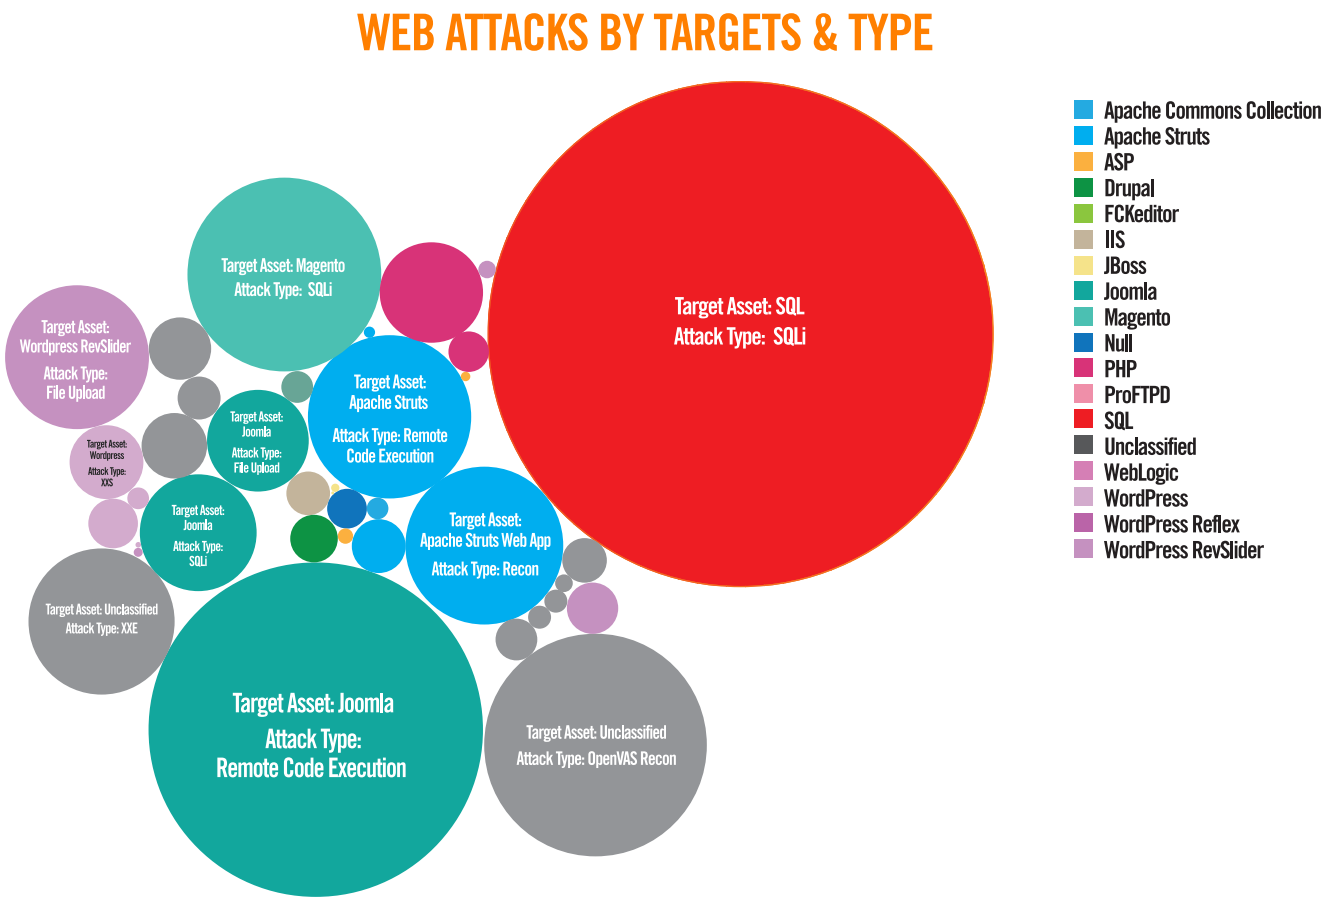

Figure 6: When web application attacks are grouped by target and type, certain popular attacker targets stand out, as we discuss below. Note the large mass of otherwise undefined SQLi issues. Alert Logic detection is able to discern these attacks at multiple points in the chain of events and within multiple layers of the app stack, so their presence is even greater than this chart would indicate.

RIPE PICKINGS: CMS AND E-COMMERCE PLATFORMS

To begin our deeper look at incident types, we turn our attention for a moment to attacks focused on CMS (content management systems) and e-commerce platforms – a rich hunting ground for attackers working in the web apps space. In this category of applications, Joomla came in with 25% of the total web application attacks sorted by incident data, followed by WordPress (10%) – heavily weighted by exploits targeting WordPress RevSlider – and Apache Struts (10%). Joomla has become a very popular framework for web applications, and our threat telemetry indicates that it has become one of the most targeted; in the second half of 2016 alone, Joomla-related web application attacks represented approximately 6% of overall observed security incidents. Increases in exploit automation and non-technical recon mechanisms like Google dorks (that is, search techniques designed to elicit information that may not easily emerge with a more casual Google search, such as structured data in spreadsheets) make targets like Joomla an easy choice for attackers.

WEB APP SECURITY ISSUES ARE NEVER "JUST" ABOUT THE WEB APPLICATION, OR EVEN THE DATA IN THE NEAREST DATABASE.

The largest known data breach to date – the 2013 Yahoo incident that affected a billion user accounts – hinged on CVE-2012-3414, a cross-site scripting flaw in WordPress that had in fact been fixed in 2012. WordPress was the CMS in use by Yahoo's developers for their team blog at the time of the breach.[7]

For a more granular look at the problems we're seeing, let's leave aside SQL and SQLi for a moment and go deeper into the content management systems and e-commerce components prominent in the chart above: Joomla, Magento, WordPress, and Apache Struts – technically speaking not a CMS or e-commerce platform but rather a foundational web application framework underpinning many CMS and e-commerce application stacks.

TOP TARGETED CMS AND E-COMMERCE STACK COMPONENTS

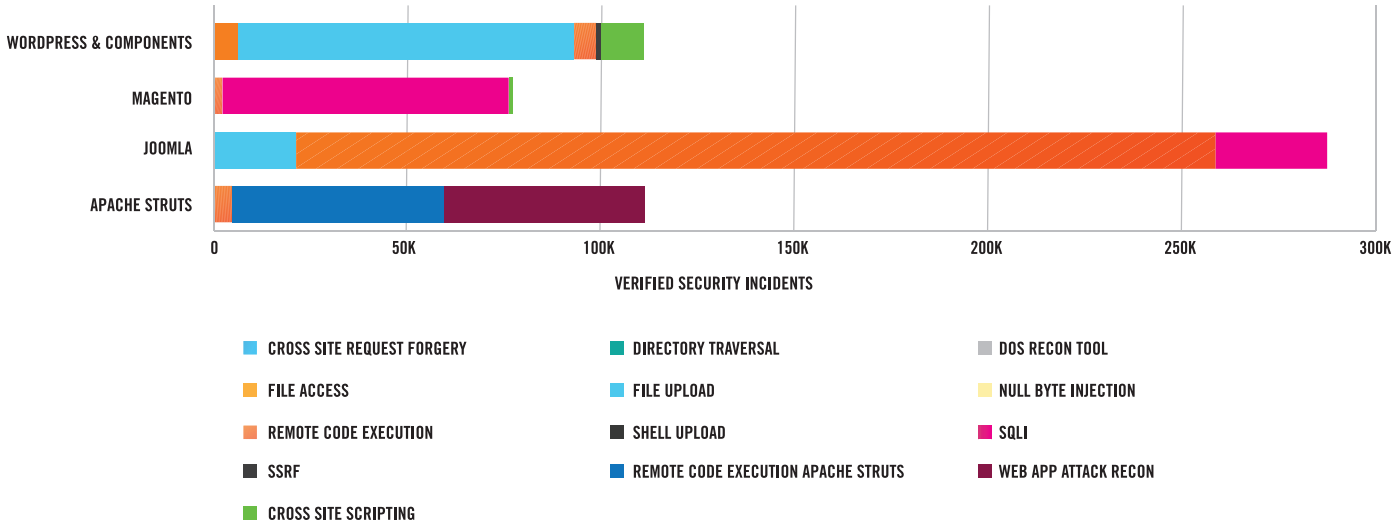

Figure 7: Targeted web application assets and exploits observed in proportion. Note that each target is predominantly associated with a particular type of attack, but that no component is only susceptible to a single type.

EXPLOITS TARGETING JOOMLA

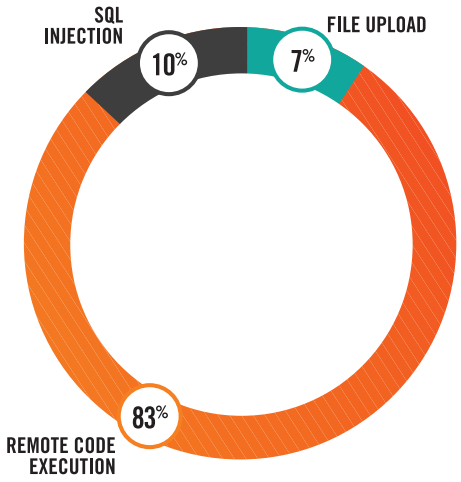

Figure 7a: Exploits targeting Joomla overwhelmingly take advantage of remote code execution vulnerabilities.

EXPLOITS TARGETING MAGENTO

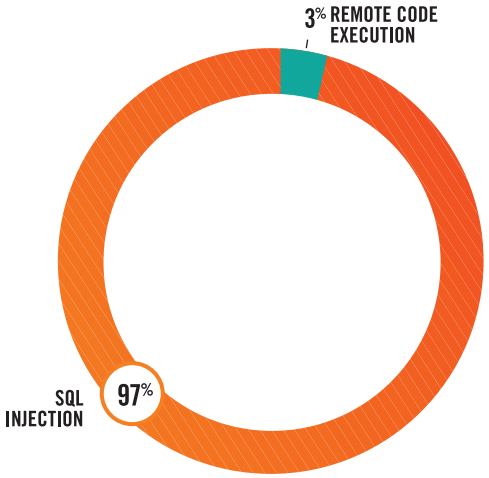

Figure 7b: Exploits observed targeting Magento e-commerce applications focused on SQL injection.

Joomla-focused attacks account for 25% of all web application attacks observed and 49% of the most targeted web app stack components represented above. Its integration with SQL databases make it a natural target for SQLi attacks, and about 10% of the incidents we saw were precisely that, but 83% of the Joomla incidents we saw involved remote code execution. This could indicate that Joomla is a sought-after entry point for lateral movement into a network.

Magento-focused attacks account for 7% of the total web application attacks observed and 13% of the top represented; of these, nearly 97% were SQLi issues that could be definitively linked to the platform. Given the high value of data possibly being stored in SQL databases within Magento application stacks, injection attacks used to exfiltrate that high value data seem appropriately represented.

EXPLOITS TARGETING WORDPRESS

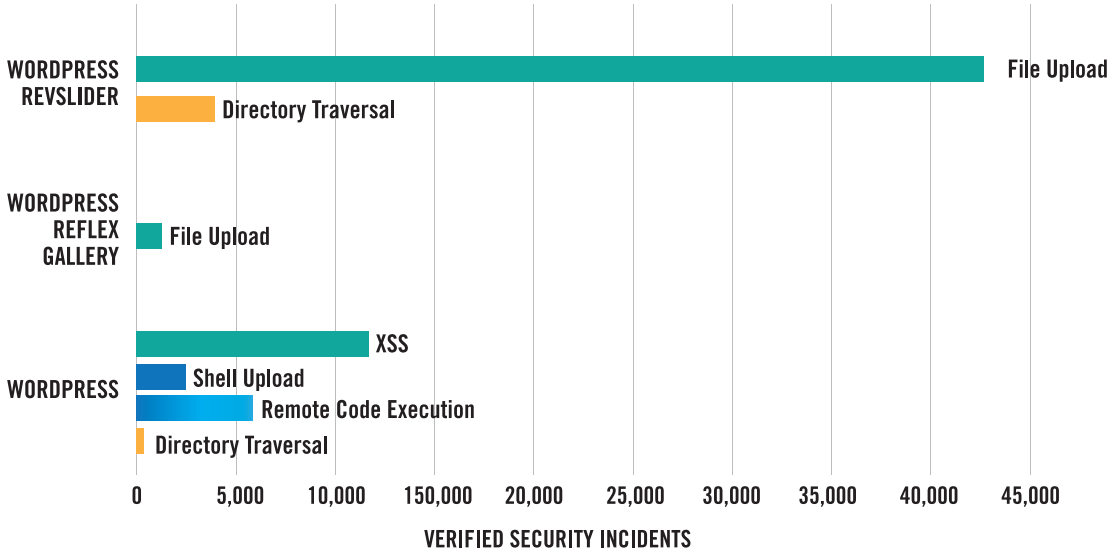

Figure 7c: WordPress remains a wildly popular content publishing tool, with a rich ecosystem of plug-ins and add-ons. However, that flexibility affects its overall security profile. In this figure, we see that exploits targeting a specific WordPress plug-in account for the lion's share of all WordPress-related issues.

WordPress-focused attacks account for around 10% of all web application attacks observed and 19% of those represented above; you'll notice that we've sorted this segment into two chart sections, assigning a brick of its own to the popular WordPress Revolution Slider plug-in, famously the vector in the massive 2014 SoakSoak attack. The other WordPress block includes incidents involving Reflex Gallery, Symposium, XML-RPC, and other WordPress or WordPress-related code. File upload issues made a strong showing for all WordPress-related incidents, accounting for 77% of their total.

As noted above, our data shows platforms such as Magento and Joomla to be specifically targeted as part of SQLi campaigns, above and beyond standard SQL-driven web application frameworks. Let's take a moment to look more closely at SQL injection, since it's clear that the concept should figure prominently in any conversation about modern data protection.

CREEPERS BY THE DOZEN: SPOTLIGHT ON MACHINE LEARNING AND SQL INJECTION ATTACKS

Certain types of web application attacks use old flaws that can best be addressed by newer, data-driven solutions. When most security reports discuss the frequency of attacks or vulnerabilities, they identify incidents atomically – that is, each instance is reported as an occurrence. The statistics in this report reflect that approach. But experienced security analysts know that in the real world, attacks rarely happen in isolation, and that multiple vulnerabilities may be exploited in a single breach.

The process of correlating incidents to form a full understanding of a common attacker or campaign is notoriously difficult. It often requires some element of guesswork or intuition by experienced security analysts. Even with the benefit of insight and experience, attempts to tie disparate attacks together are imprecise, and even the best human analysts can sometimes only piece together weeks- or months-long campaigns in hindsight. This work helps the next victims – but as the saying goes, the pioneers tend to end up with arrows in their backs.

Of all the attack types we analyze in the Alert Logic SOC (Security Operations Center), this problem is most pronounced with SQLi (SQL injection) attacks. In the 18 months of data our report covers, SQLi accounts for 634,282 incidents, or 55% all observed attacks. SQLi is a well-understood attack technique – hacker publications such as Phrack were discussing it all the way back in 1998 – and the rise of the data-powered Web has caused its popularity among attackers to explode. Each incident involving SQLi may consist of hundreds of events, the vast majority of those being unsuccessful execution attempts. Even more confusingly, as the illustration

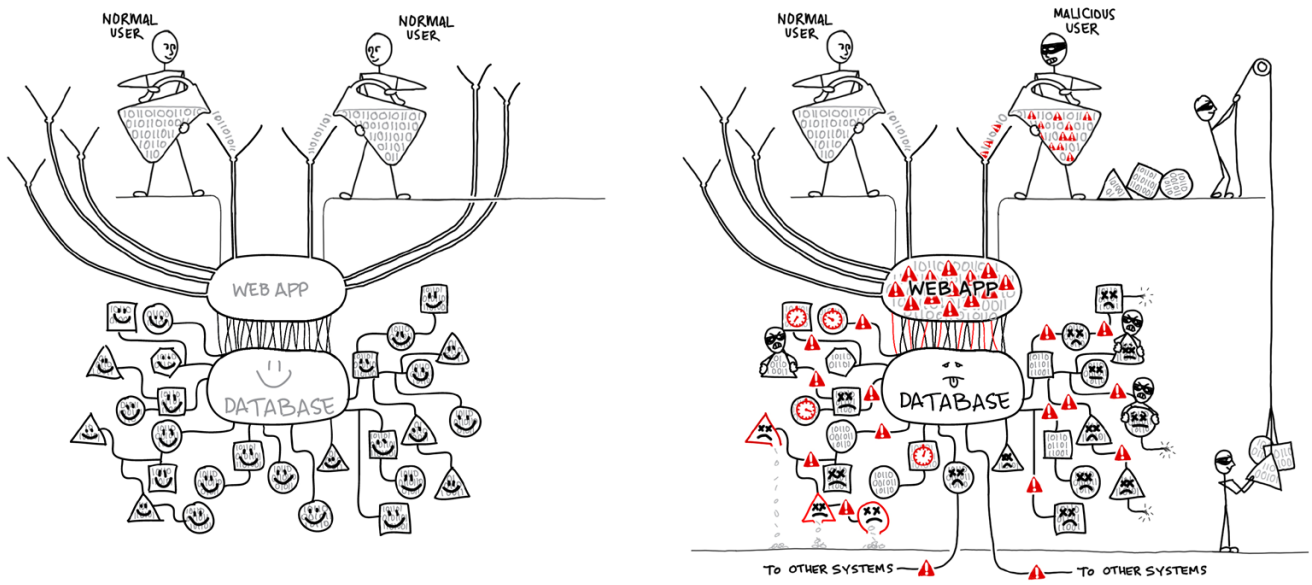

On the left, well-behaved users use a web app to interact with a database. On the right, a malicious user pours in malicious inputs – and the web app, unprepared to sanitize or reject such inputs, passes the code along. The database is infected, with data inaccessible, corrupted, deleted, or exfiltrated.

above shows, the process of code injection takes advantage of certain cloud-crucial advances in processing, such as those that improve processing speed. Identifying real-world attack campaigns in such a dataset is notoriously difficult – at least if we rely exclusively on humans for analysis.

Enter machine learning – a marvelous means of tackling a large, ambiguous, data-intensive attack set such as SQL injection. Over the past year, our researchers have developed new ways to deploy machine learning tactics against SQLi's titanic datasets, and while we have less than a year of results to share in this report, we're including our first crop of findings as a preview of what you should expect from Cloud Security Reports in the future.

Over the nine-month period during which our machine learning effort came to life, 44.2% of successful SQLi attacks resulted in disclosure of information about the database (for instance, the version of the database in use, or information on the schema), but showed no identifiable evidence of progression beyond that point. Another 38.5% of attacks got a bit further and were able to disclose more information – tables and/or field names, row and column limits, and such – and showed the potential to inject commands, but couldn't manage to breach actual data. And a small but toxic 17.3% made it all the way to the goal, exfiltrating row data at ease and breaching (or dumping) data.

Most significantly, we identified approximately 231 attacks, or about three-quarters of those in the smallest and most toxic group of incursions, in which malicious SQL injection was deployed with a high degree of complexity and sophistication – superior knowledge of (and skill at) breaching database architecture and bypassing well-understood IDS detection methods. While this may seem like a small part of the larger picture, it means that 8-10% of the customers we monitored were targeted by actors with better-than-average levels of skill and determination, which is notable. These actors used various forms of obfuscation and personally created sophisticated SQL injection attacks more complex than common tooling – such as SQLMap, Burp Collaborator Suite or Havij – support.

Just over half (53%) of these 231 incidents were detected mainly by use of our more traditional detection and analysis methods. In those cases, machine-learning techniques allowed us to better understand attack progression and to provide meaningful context for the attack the SOC had identified. The remaining 47% of these incidents were detectable only with the use of machine learning.

ET TU, BRUTE FORCE?

Brute force has been one of the top threat vectors across all industry verticals, geographies, and environments for quite some time. With brute force attacks, we see that the exploitation of vulnerabilities considered "path of least resistance" continues to increase, highlighting the need for organizations to continue implementing more stringent and tightly controlled identity and access management (IAM) policies. (It also reminds us that attackers are essentially not biased toward innovation; the tool that already works, or even mostly works, is preferable to the tool that doesn't exist yet.)

52%

OF ALL BRUTE FORCE ATTACKS WE SEE TARGET WINDOWS PLATFORMS

Brute force attacks are more prevalent in on-premises environments, representing about 12% of the total incidents we noted during the timeframe of this report. The range of brute force attacks is wide, given that a persistent or professional attacker is likely looking for corporate secrets and thus may try for access to email servers, login credentials, or intelligence. When digging into the data, most of the suspicious brute force activity we detected was linked to account creation and security group modification, and mostly in pursuit of escalated system privileges and exfiltration of data.

Given the telemetry provided, it is clear that Windows remains the most targeted platform across Alert Logic's threat data. Approximately 52% of all brute force attacks we see target Windows platforms in the data center, representing a key trend in how attackers target publicly available

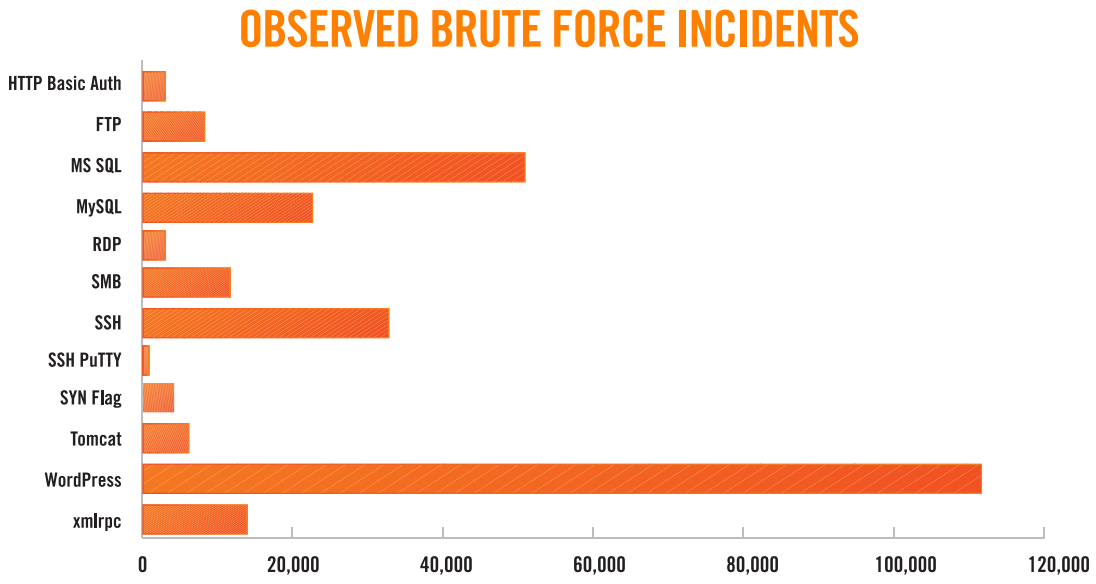

Figure 8: WordPress-related attacks lead the brute force category by a large margin.

services across organizations. However, with many systems and platforms maintaining administrative access via SSH, it's clear that SSH brute force attacks against customer networks are on the rise. Seven percent of brute force attacks we noted target SSH.

WordPress, currently the Internet's most popular web application framework (27% of all sites use WordPress, and not merely for blogging), gives us an excellent example for walking through a brute force attack. It also spotlights the role that content management systems play in attack progressions – they're rarely the ultimate target, but they're crucial stepping stones on the way there. WordPress, now deployed far beyond its humble blogging origins, is notorious for being susceptible to a smorgasbord of threats, thanks to certain questionably constructed plugins and themes offered on the community download sites. Naturally this lack of good plug-in behavior has invited increased threat activity, with WordPress-related attacks representing approximately 17% of overall attacks across Alert Logic's threat data.

To execute one popular brute force attack against WordPress, the attacker sends multiple POST XML requests with one or multiple sets of credentials, in the hope of finding a login to the WordPress website. If the attacker receives an "Incorrect User" with a 403 ("Forbidden") code returned in the response from WordPress, the brute force attack fails. However, if WordPress responds with 'isAdmin':True in the response, the attacker has succeeded in finding a way into the web site.

Other common brute force incidents focus on specific protocols such as SSH, FTP, SMB, and RDP. These attacks take place in various ways; the majority are generic password-guessing attacks targeting poor passwords (e.g. 123456, qwerty, or password) and typically seeking accounts such as root or admin that exist in every organization. Alternately, an attacker might seek accounts that might exist on a system. based on the technology (eg., Oracle, SQL, or Cisco) known to be in use. In either case it should be noted that the definition of a "poor" password has a lot of elasticity; the examples we give above are particularly awful, but in the modern era it's trivially easy for an

A brute force attack is slightly misnamed – less a goon squad putting a boot to the door, more a constant swarm-like effort to eventually achieve the desired goal. Modern technology and processing advances means that brute-force attacks can find success in far less time than they might have previously.

attacker to automate brute force attempts that can run through every word in the dictionary, most popular given names, and all the obvious numeric combinations under a certain length in next to no time.

The brute force attempts we saw were detected mainly through our system log and network monitoring technologies, which incorporate efficient security content able to detect lateral moves and account modifications. These attacks, and the efficacy of these detections, show the importance of securing all layers of your infrastructure using a solid, in-depth security strategy… including cleverly constructed passwords, changed regularly.

SERVER-SIDE RANSOMWARE: THE CONTINUING SAGA

Malware represents a diverse constellation of events, even though incidents serious enough to trigger Alert Logic notification represent only 2 percent of our total incidents in this report. Despite this, the malware class of threats remains one of the most interesting to customers, given that the goal of threat detection is to ultimately prevent or stop malicious applications or activity from happening within the infrastructure.

The interest is sadly justified, as our coverage period saw the rise of an infection type that, if not entirely new under the sun, promises to wreak a very special kind of havoc: server-side ransomware.

Server-side ransomware is a painful exception to the rule that says hackers are fairly slow to innovate if they've got a really effective tool in hand – since, after all, there continues to be no shortage of users that will click on suspicious links or email attachments. But users occasionally log off, and their importance is usually relatively minimal compared to that of servers – large, always-on, often-unpatched servers full of tasty and valuable data. The target was simply too tempting for attackers to resist.

And resist it they did not. Using – once again – a chain of vulnerabilities and exploits, attackers in March 2016 unleashed SamSam and Maktub. They targeted mainly healthcare installations (which kept the overall incident numbers low in our dataset), hopscotching from system entry to lateral movement to malware installation to offline encryption and ransom demands.

In a way, demands for Bitcoin are one of the few ties between this new kind of malware and what we traditionally think of as ransomware. Maktub and SamSam attackers need at least some degree of knowledge about the target to establish a foothold, while traditional end-point focused malware relies mainly on volume, operating system prevalence, and bad user habits to do its work. The new breed of server-side malware is also able to handle certain tasks, such as encryption, without phoning home to a command-and-control server. This poses interesting new challenges for detection and forensics.

AS THIS REPORT WAS ENTERING PRODUCTION, THE WANNACRY RANSOMWARE ATTACK WAS BURSTING INTO THE HEADLINES, WITH OVER 230,000 USERS AND 10,000 ORGANIZATIONS IN 150 COUNTRIES AFFECTED IN THE FIRST FEW DAYS. THIS RANSOMWARE TARGETS WINDOWS INSTALLATIONS AND USES A VULNERABILITY DISCLOSED AND PATCHED IN EARLY MARCH. VARIANTS ARE ALREADY TURNING UP ONLINE. WELCOME TO THE FUTURE!

Finally, the business model for the new breed of server-side ransomware is still apparently being worked out; while traditional ransomware has been around long enough for attackers to have an exquisite sense of what the market will bear, the March 2016 attackers appeared to still be working out the kinks in their demands[8]. Unfortunately, the odds are they will have enough time to figure out exactly how much their unpatched and unfortunate targets are likely to pay to regain access to their data.

SECURITY FAIL

A MESSAGE YOU CAN BUG

A wise person once advised his listeners to make everything as simple as it can be, but not more so. That's a lesson the proprietors of CloudPets learned the hard way in early 2017.

A CloudPet is a stuffed animal connected to an online account, to which anyone with the right sign-in information could send audio messages that would play through the toy to, presumably, the child in possession of it. All very snug and cuddly, and the tagline was "a message you can hug" – what's not to like?

However, the security underpinnings turned out to be less than likeable. A MongoDB behind the scenes held names and hashed passwords – but parent company Spiral Toys did not enforce any minimum strength for those passwords, so even a single letter would do. In turn, the database itself was online but required no authentication for access; likewise, the actual audio messages were stored in an Amazon S3 bucket with no authentication. The MongoDB contained reference file paths pointing to the S3 buckets, so in theory an attacker could easily move between the two, deleting files or pointing to other files at will.

Though it's believed that the databases were the subject of ransomware demands at least twice, it's unclear whether voice recordings were in fact stolen or redirected from individual accounts – the company strenuously denied it, though researchers easily demonstrated proof of concept. The creepier situation arose when recordings were…added. At the other end of the chain, researchers also found hardware security to be lacking, with no pairing protections on the toy's Bluetooth – allowing anyone within range to pair with the toy, record a message, and play it back.

Ransomware, breached data, reputational damage, flashbacks to that Twilight Zone episode with "My name is Talking Tina and I don't LIKE you!" – it's a security story that lacks only frosty security community relations and a Congressional inquiry to be complete, and actually it lacks neither. Researchers attempting to privately disclose issues to the company as early as October 2016 were unable to make contact, while the company claimed it knew nothing of the situation until late February. Within days, Sen. Bill Nelson (D-FL) was reaching out to the company for information on how the breach happened [10], what has been done to rectify the situation, and further information on how the company is complying with COPPA (Children's Online Privacy Protection Act) regulations. By late April, indications were that the company was in the process of going out of business[11] – with no word to customers on what that might entail, or what the future held for their data (or their toys); however, they could still be purchased through Amazon (or through Amazon-using retailers).

> *"The MongoDB contained reference file paths pointing to the S3 buckets"*

What might CloudPets have done differently before the breach? Clearly the company understood that its audience might not be comfortable with strong passwords, but that's no reason for hilariously lax password-creation rules – and relying on hashes (almost certainly introduced by the libraries used by the developers, rather than by some conscious choice on the company's part) as the sole line of defense for not just individual accounts but the entire multi-cloud was fatal. Best practices would have seen the company encrypting all the user information – and, of course, ensuring that all databases involved used strong authentication.

CASE STUDY

SECURITY WIN

OH, CANADA!

Many security stories – most, perhaps – are tales of woe, of things that went wrong or weirdly. It's important to hear and learn from examples of good response, too, so let's look at a recent success: The Canada Revenue Agency's smart, rapid response to the Apache Struts 2 bug in March 2017.

To recap the incident itself, the attackers uploaded a maliciously constructed file that included an invalid Content-Type, Content-Disposition, or Content-Length value[12]. The Jakarta Multipart parser used by the Struts 2 framework would respond with an exception, which the attackers could then use to achieve remote code execution on the target machine. After that, the activity varied, but widespread reports included instances of dropped payloads, disabled security tools, and persistent attacks. The issue affected Struts 2.3.x through 2.3.31, and Struts 2.5.x through 2.5.10. The Apache bulletin number was S2-045[13] and the eventual MITRE identifier was CVE-2017-5638[14].

The Struts 2 vulnerability was publicly disclosed on March 6 – a Monday – and was weaponized and active within 24 hours[14]. On Wednesday, Canadian officials detected evidence of an attack; on Friday at midnight, the agency pulled its site offline briefly as a precaution, restoring everything but then pulling its digital services back offline around 2 PM that day; all was back to normal on Sunday around 5 PM. The agency has stated that no taxpayer information was breached.

Unpacking the events, we see an admirable balance of attentiveness, caution, forthrightness, and overtime/weekend hours. First, the agency didn't wait for indications that they were breached; they stayed on top of an active security situation and didn't wait to feel the burn for themselves. And they kept checking as the situation developed; the check in the wee hours of Friday morning did not apparently reveal the potential issue with the digital property, but whoever checked the "ok" box in the morning did not hesitate to uncheck it in the afternoon as events dictated.

Second, someone was empowered to make the tough decision to pull the site offline not once but twice – first to check everything over, and then again when a property not previously known to be susceptible was found to be so. This choice may be bolder than it appears on first blush; the Canadian income-tax filing deadline on April 30, but tax preparers are already hard at work in early March, and there were scattered reports of shutdown-related inconvenience to preparers and taxpayers. Balancing system requirements for integrity and availability is never simple, but CRA found a way[15].

Government officials credited the nation's Shared Services Canada system with making the process work. The system coordinates federal IT services as one would do in a large enterprise, rather than allowing agencies to operate in silo-fashion. In a press briefing the following Monday, the COO of Shared Services Canada noted that another Canadian agency was also found to be vulnerable and was handled in much the same fashion and on the same timeframe[16]. Throughout, the agency communicated shutdowns and other pertinent information to the public – not too much information, but enough that it was clear that matters were in hand.

And the overtime hours? The price of doing business, one supposes – but perhaps an easy calculation to make when weighed against the potential for not only data breaches but delays in tax revenue. The CRA reported that not only was the effort successful, but timely action meant that they expected no delays in processing 2016 returns – and had no plans to offer filing extensions to taxpayers. Smooth response may not have resulted in a tax extension for the citizens, but the payout in safety and security was worth even more.

THE MORE THINGS CHANGE: THE VERTICAL VIEW

What do clothing stores, hospitals, factory floors, and banks have in common? Quite a lot, as it happens – their back offices have many of the same job roles and rely on many of the same tools, a truth that was richly evident in the data examined for this report. It's one of those findings in the category of "boring but important" – and a very important reminder that the high level of system similarity and interdependency among workplaces is a contributing factor in attacker mobility and ease of access these days.

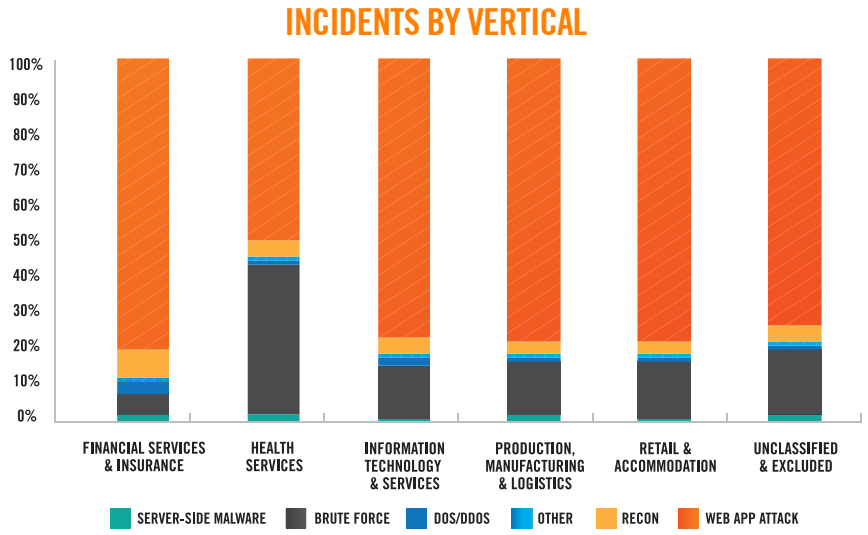

Figure 9: When sorted by affected vertical, it's easy to see that attackers find certain kinds of exploits widely useful.

Since attackers don't hesitate to learn from each other and apply successful tools, tactics, and procedures to get the job done, defenders must do the same. The following charts show both vast similarities (as you would expect at this point in the report, SQLi attacks absolutely dominate the dataset for all sectors) and, occasionally, some intriguing differences.

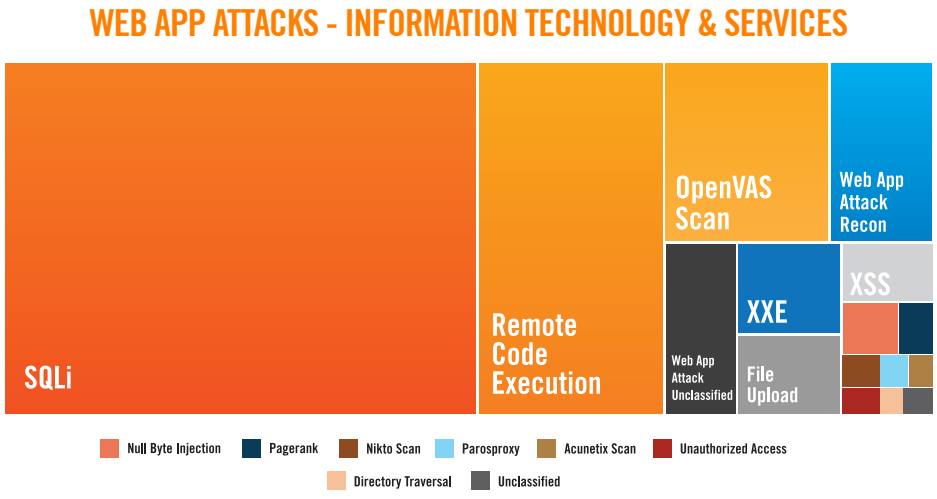

Figure 10: The web application attack landscape for the Information Technology and Services sector.

Our first chart, showing the Information Technology and Services sector, sets a soon-to-be familiar hierarchy, with SQLi incidents representing the lion's share of mayhem, followed by remote code execution-related activities and events related to Open Vulnerability Assessment System scans.

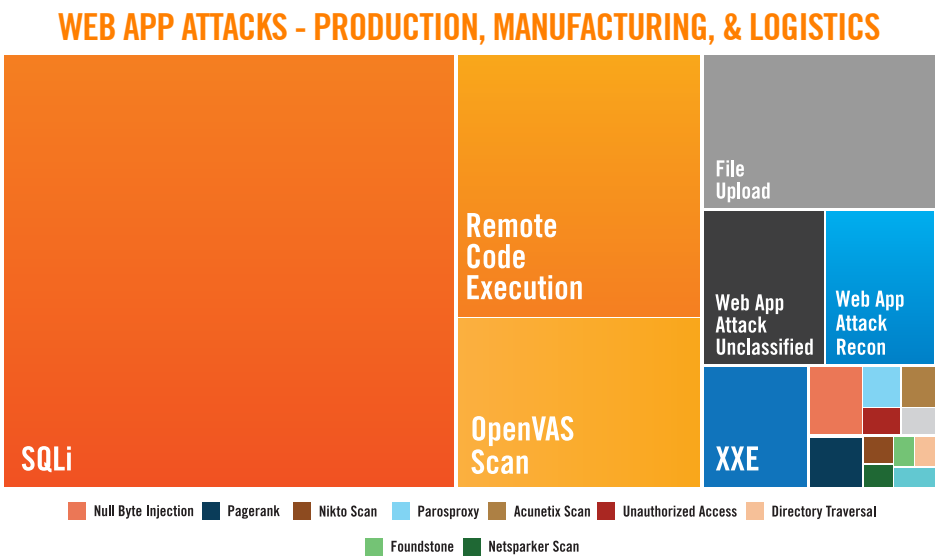

Figure 11: The web application attack landscape for the Production, Manufacturing, and Logistics sectors.

The SQLi-RCE-OpenVAS triumvirate reigns in the Production, Manufacturing, and Logistics sector as well, though problems related to file uploading (many related to specific WordPress issues) make a strong showing as well. Cross-site scripting, a notable factor in the Information Technology and Services table, is less common in this sector.

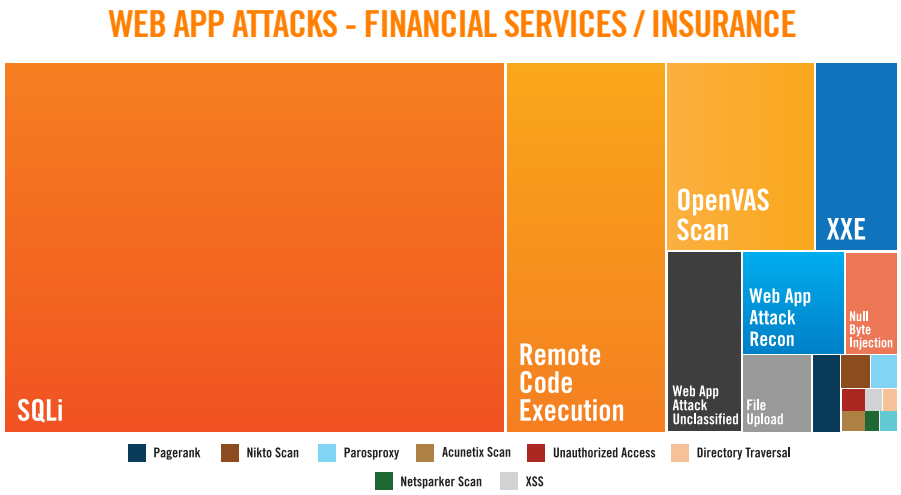

Figure 12: The web application attack landscape for the Financial Services / Insurance sector.

SQL injection is more significant than ever to the Financial Services/Insurance sector, with remote code execution also making a particularly strong showing and OpenVAS-related incidents the third most common finding. We also see XXE (XML External Entity) vulnerabilities, in which an unreliable source attacks an application that parses XML inputs, making some noise in this sector.

WEB APP ATTACKS - RETAIL AND ACCOMMODATIONS

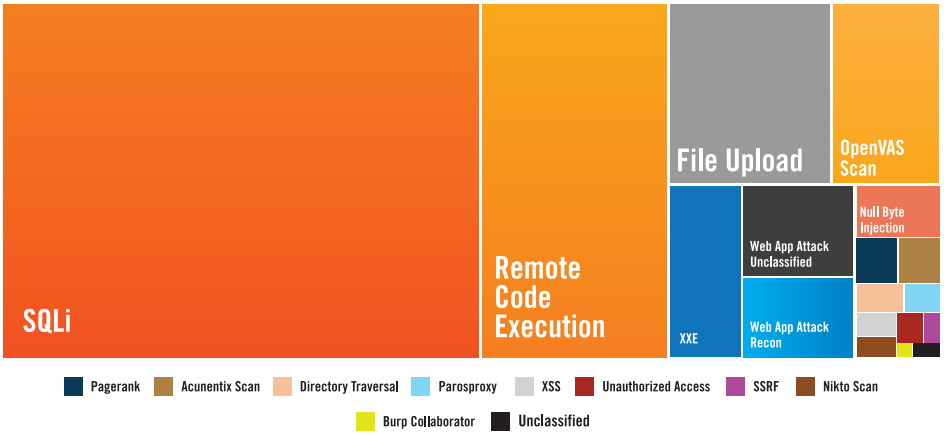

Figure 13: The web application attack landscape for the Retail and Accommodation sector.

We have an upset of sorts in the Retail and Accommodation sector, with File Upload issues prevalent enough to knock OpenVAS out of its usual third-place perch. SQLi and remote code execution are, however, not to be denied their usual spots at the head of the table.

WEB APP ATTACKS - HEALTH SERVICES

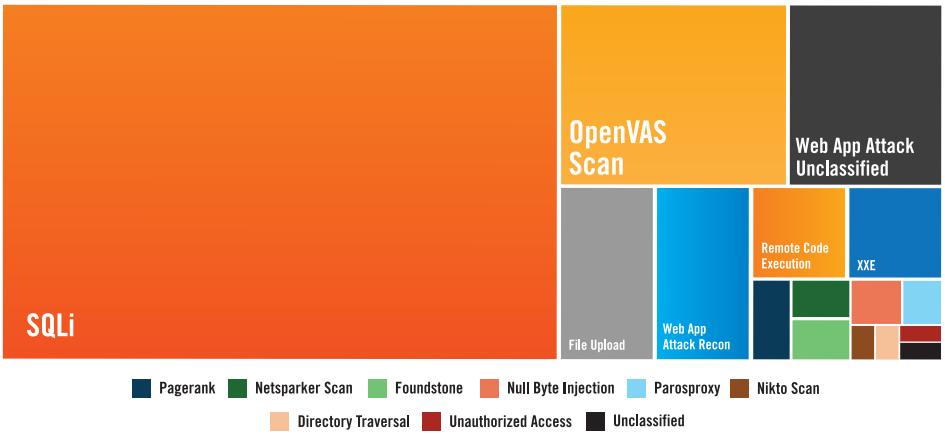

Figure 14: The web application attack landscape for the Health Services sector.

Our fifth sector, Health Services, has more surprises in store – not in the utter dominance of SQLi issues, perhaps, but in the relative lack of traction for remote code execution in this part of the dataset. OpenVAS instead returns to play first runner-up.

With so much uniformity, it's fascinating to see how sector-focused individual attacks can profoundly affect a single sector's numbers. For instance, the Health Services sector had an interesting summer in 2016.

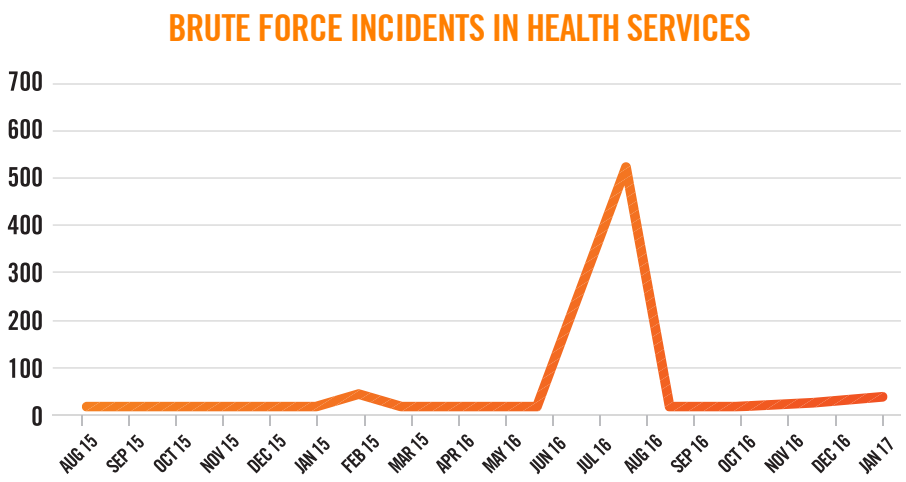

Figure 15: A midyear spike in incidents targeting health services is ascribable to attacks against three specific Alert Logic customer targets.

This visible jump in the rate of brute force attacks against enterprises in the Health Services sector reflected an increase in attacks in early summer against just three specific Alert Logic customer targets. When those attacks ceased, the proportion of brute force attacks to others in the sector resumed normal levels. Careful analyst and machine learning assessment of the data as it flowed in helped us to pinpoint the problem – not a pandemic, just a nagging seasonal itch.

The consistency in attack type across verticals is more remarkable when you know that network types are not evenly distributed by vertical in our data; as one would expect, some verticals are far farther along in cloud adoption than others.

CONCLUSION

And now? With attack surfaces changing and so much more likely to come, what can you do to optimize your enterprise's security? The advice below, adapted from that offered by US-CERT[9], should hang on the wall of every Information Security office, regardless of whether the business uses public, hybrid, or on-premises computing power – it's a baseline for good security practice wherever your data resides.

To prevent targeted attacks:

- First, rely as much as possible on application whitelisting. Blocking access to unknown programs can keep malicious applications from gaining access to the network and its assets. And never be afraid to take a hard risk-assessment look at the value an app adds versus the risk to which it exposes you.

- Second, as we've seen, applications are a primary target, and vulnerabilities within applications are the easiest means of exploiting those. Patching matters. We still see regular evidence of attack attempts by such venerable vermin as Conficker (est. 2008) and the Slammer worm (est. 2003). Continually assess your attack surface for vulnerabilities and configuration exposures. Understand your own patching process and, when patches are available, make it a priority to evaluate and deploy them. If your providers don't provide notifications and clear communications about security issues, insist that they improve their customer service.

- Finally, remember that users, or at least their access to your system, are always a potential problem. Restricting administrative and access privileges based on current user duties can prevent malware or other types of attacks from spreading. Privileges for both applications and operating systems should be kept up to date.

AND BEYOND BASICS:

Look at how your organization monitors emerging threats, and form good relationships with the vendors and services that make it their business to keep an eye on the landscape. Know your applications, how your portfolio impacts your attack surface, and keep a close eye on vulnerability disclosures involving them. This goes for your operating systems and even your hardware.

Don't dismiss the risk posed by older vulnerabilities and exploits. As we discussed earlier, attackers don't usually worry about showing off the new hotness in malicious attacks; they use whatever works most efficiently and reliably.

Likewise, know that people don't change. Humans inside your organization can still cause a formidable amount of trouble. Diligent employee training will cut down on inadvertent mishaps and that's good, but event logging is necessary. Your logging policies should be as clear and well thought out as your data retention policies…which you should also have in place, and reviewed with the same executive diligence.

Finally, as attack surfaces change, be ready to make clear choices about where your enterprise belongs on the continuum from on-premises hosting, to hybrid or managed hosting, to a full public cloud solution. Know your infrastructure and IT staffing plans, and keep abreast of how your business-critical applications and services may be threatened. There is no single sit-back-and-relax solution; the day to day flow of security management will change depending on where you choose to place yourself on the continuum. Ultimately, the responsibility to choose good security partners and to protect your employees, customers, and IP are your own.

APPENDIX A: THE DATA

SECTION 1: AN OVERVIEW OF ATTACK TYPES AND TARGETS

SUM OF ALL ANALYZED SECURITY INCIDENTS BY TYPE

Server-side Ransomware	Brute Force	DoS/DDoS	Other	Recon	Vendor Scan	Web App Attack*	Grand Total
27,307	271,593	24,815	17,444	79,193	501,313	1,286,130	2,207,795
1%	12%	1%	1%	4%	23%	58%	

SUM OF SECURITY INCIDENTS BY TYPE, LESS VENDOR SCAN

Server-side Ransomware	Brute Force	DoS/DDoS	Other	Recon	Web App Attack*	Grand Total
27,307	271,593	24,815	17,444	79,193	1,286,130	1,706,482
2%	16%	1%	1%	5%	75%	

**Web App Attack combines Web App Attack Recon security incidents with Web App Attack security incidents*

AVERAGE INCIDENTS PER CUSTOMER BY ENVIRONMENT

	Incident Count	Average Incidents per Customer
Public Cloud	133,701	405
Hosted Private Cloud	1,589,415	684
Hybrid	28,328	977
On-Premises	348,315	612

TOP OBSERVED INCIDENTS - PUBLIC CLOUD

Incident Type	Incidents	Average Incidents per Customer
SQLi Recon	15,286	1,092
Joomla Web App Attack	12,324	76
SQL Injection	4,701	44
Apache Struts Recon	4,700	23
Magento SQL Injection	3,786	64
WordPress Brute Force	3,400	71

TOP OBSERVED INCIDENTS - HYBRID

Incident Type	Incidents	Average Incidents per Customer
SQLi Recon	4,953	225
Joomla Web App Attack	3,315	158
SQL Injection	2,501	125
WordPress Brute Force	2,400	200
Magento SQL Injection	1,424	158
Apache Struts Recon	669	28

TOP OBSERVED INCIDENTS - HOSTED PRIVATE CLOUD

Incident Type	Incidents	Average Incidents per Customer
SQLi Recon	290,334	194
Joomla Web App Attack	199,906	126
SQL Injection	99,192	93
Magento SQL Injection	60,877	75
MS SQL Brute Force	31,471	213
SIPVicious Scan	21,704	261
WordPress Brute Force	20,711	149
Apache Struts Recon	12,147	24

TOP OBSERVED INCIDENTS - ON-PREMISES

Incident Type	Incidents	Average Incidents per Customer
SQLi Recon	51,296	209
Joomla Web App Attack	24,538	116
SQL Injection	22,685	98
WordPress Brute Force	16,177	88
SIPVicious Scan	8,172	101
Vega Scan	7,610	29

SECTION 2: SECURITY INCIDENTS BY TYPE

WEB APP ATTACK INCIDENTS BY TARGET ASSET

Type	Observed %
Apache Commons Collections	0.13%
Apache Struts	10.11%
ASP	0.04%
Drupal	0.43%
Elasticsearch	0.03%
FCKeditor	0.00%
IIS	0.32%
JBoss	0.02%
Joomla	26.11%
Magento	6.98%
PHP	2.15%
ProFTPD	0.02%
SQL	47.74%
VMWare VCenter	0.00%
WebLogic	0.01%
WordPress	1.61%
WordPress Reflex Gallery	0.05%
WordPress RevSlider	4.19%
WordPress Symposium plugin	0.02%
WordPress XMLRPC	0.04%
Total	100.00%

BRUTE FORCE INCIDENTS BY TARGET

Brute Force Vector	Observed Incidents	% of Total
HTTP Basic Auth	5,122	2%
FTP	8,197	3%
MS SQL	50,807	19%
MySQL	22,551	8%
RDP	4,382	2%
SMB	11,617	4%
SSH	32,705	12%
SSH PuTTY	537	0%
SYN flag	4,138	2%
Tomcat	6,097	2%
WordPress	111,541	41%
XML-RPC	13,898	5%
Total	271,592	100%

WEB APP ATTACKS TARGETING TOP CMS AND ECOMMERCE ASSETS

	Cross Site Request Forgery	Directory Traversal	DoS recon Tool	File Access	File Upload	Remote Code Execution	Shell Upload	SQLi	Recon	XSS
Apache Struts	0	0	0	0	0	60,129	0	0	51,236	0
Drupal	0	0	0	0	0	0	0	4,691	0	0
Joomla	0	0	0	0	21,604	237,367	0	28,579	0	0
Magento	0	0	0	0	0	2,327	0	74,590	0	5
WordPress	13	0	70	2	0	5,418	874	0	0	11,210
WordPress Reflex Gallery	0	0	0	0	505	0	0	0	0	0
WordPress RevSlider	0	3,446	0	0	42,594	0	0	0	0	0

TOOLS DRIVING MALICIOUS RECON INCIDENTS

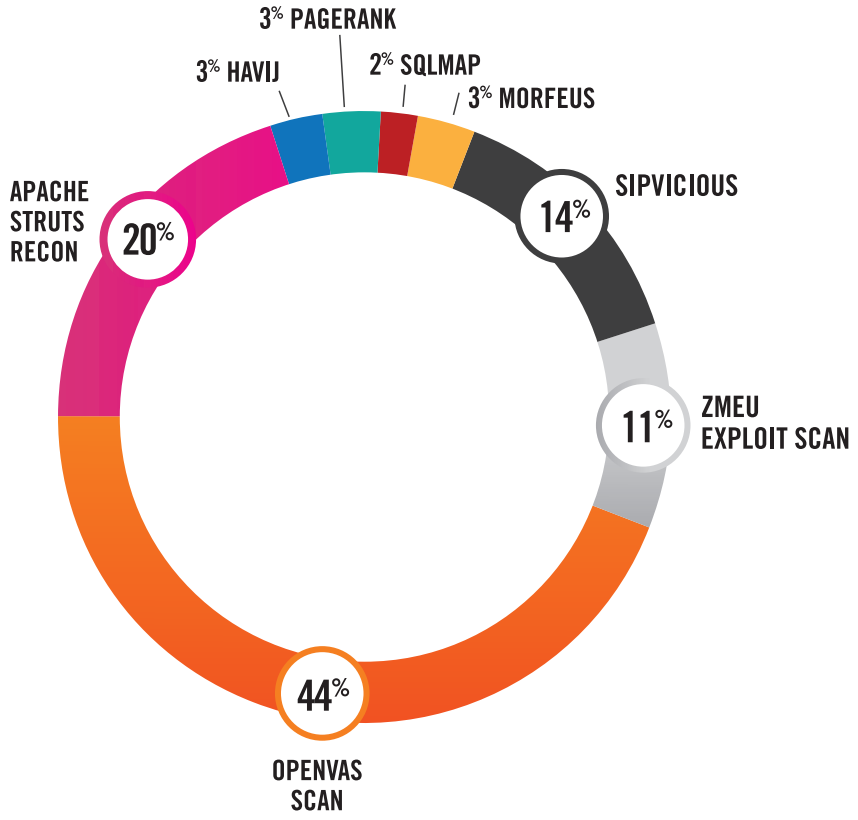

OBSERVED SECURITY INCIDENTS BY ENVIRONMENT AND INDUSTRY

	Financial Services & Insurance	Health Services	Information Technology & Services	Production, Manufacturing, & Logistics	Retail & Accom.	Other	Grand Total	Infrastructure %'s
Public Cloud	8,041	3,446	54,018	20,852	14,326	33,018	133,701	6%
Hosted Private Cloud	102,949	45,832	541,388	248,699	199,821	450,726	1,589,415	72%
Hybrid	334	580	8,148	799	26	18,441	28,328	1%
On-Premises	17,104	5,556	83,275	67,102	51,194	124,084	348,315	16%
Unclassified	0	0	0	0	0	108,036	108,036	5%
Grand Total	128,428	55,414	686,829	337,452	265,367	734,305	2,207,795	100%
Industry %'s	6%	3%	31%	15%	12%	33%		

OBSERVED WEB APP ATTACK TYPES BY VOLUME AND INDUSTRY

HEALTH SERVICES

Web App Attack Subtypes	Security Incident Count	% of Industry Total
Acunetix Scan	5	0.02%
Authentication bypass	1	0.00%
Burp Collaborator	1	0.00%
Directory Traversal	44	0.21%
DoS recon tool	1	0.00%
File Access	4	0.02%
File Upload	1,098	5.31%
Foundstone Scan	119	0.58%
Injection	126	0.61%
Netsparker Scan	46	0.22%
Nikto Scan	105	0.51%
Null Byte Injection	2,332	11.28%
OpenVAS Scan	173	0.84%
Pagerank	82	0.40%
Parosproxy	630	3.05%
Remote Code Execution	12	0.06%
Shell Upload	12,926	62.52%
SQLi	39	0.19%
SSRF	1,294	6.26%
Unauthorized Access	1,045	5.05%
Web App Attack *unclassified*	555	2.68%
Web App Attack Recon	5	0.02%
XSS	97	0.13%
XXE	3,627	4.74%

RETAIL & ACCOMMODATION

Web App Attack Subtypes	Security Incident Count	% of Industry Total
Acunetix Scan	996	0.61%
Authentication bypass	7	0.00%
Burp Collaborator	143	0.09%
Cross Site Request Forgery	1	0.00%
Directory Traversal	644	0.39%
DoS recon tool	2	0.00%
File Access	7	0.00%
File Upload	15,812	9.65%
Foundstone Scan	17	0.01%
Injection	3	0.00%
Netsparker Scan	59	0.04%
Nikto Scan	431	0.26%
Null Byte Injection	1,964	1.20%
OpenVAS Scan	10,722	6.55%
Pagerank	999	0.61%
Parosproxy	578	0.35%
Remote Code Execution	30,639	18.70%
Shell Upload	88	0.05%
SQLi	82,259	50.21%
SSRF	315	0.19%
Unauthorized Access	402	0.25%
Web App Attack *unclassified*	5,284	3.23%
Web App Attack Recon	4,875	2.98%
XXE	6,752	4.12%
XSS	506	0.31%

FINANCIAL SERVICES & INSURANCE

Web App Attack Subtypes	Security Incident Count	% of Industry Total
Acunetix Scan	115	0.15%
Authentication bypass	3	0.00%
Burp Collaborator	17	0.02%
Directory Traversal	83	0.11%
DoS recon tool	4	0.01%
File Access	12	0.02%
File Upload	1,272	1.66%
Foundstone Scan	12	0.02%
Injection	2	0.00%
Netsparker Scan	66	0.09%
Nikto Scan	238	0.31%
Null Byte Injection	1,305	1.71%
OpenVAS Scan	6,495	8.49%
Pagerank	502	0.66%
Parosproxy	219	0.29%
Remote Code Execution	13,746	17.98%
Shell Upload	13	0.02%
SQLi	42,648	55.78%
SSRF	2	0.00%
Unauthorized Access	124	0.16%
Web App Attack *unclassified*	3,148	4.12%
Web App Attack Recon	2,472	3.23%
XSS	97	0.13%
XXE	3,627	4.74%

INFORMATION TECHNOLOGY & SERVICES

Web App Attack Subtypes	Security Incident Count	% of Industry Total
Acunetix Scan	916	0.23%
Authentication bypass	12	0.00%
Burp Collaborator	72	0.02%
Cross Site Request Forgery	8	0.00%
Cross Site Scripting	7	0.00%
Directory Traversal	742	0.19%
DoS recon tool	39	0.01%
File Access	42	0.01%
File Upload	10,125	2.55%
Foundstone Scan	147	0.04%
Heartbleed	3	0.00%
Injection	8	0.00%
Netsparker Scan	292	0.07%
Nikto Scan	1,423	0.36%
Null Byte Injection	3,632	0.92%
OpenVAS Scan	36,561	9.21%
Pagerank	2,060	0.52%
Parosproxy	1,140	0.29%
Remote Code Execution	78,594	19.80%
Shell Upload	159	0.04%
SQLi	202,855	51.11%
SSRF	50	0.01%
Unauthorized Access	1,234	0.31%
Web App Attack *unclassified*	15,180	3.82%
Web App Attack Recon	21,963	5.53%
XXE	11,996	3.02%
XSS	6,933	1.75%

PRODUCTION, MANUFACTURING, & LOGISTICS

Web App Attack Subtypes	Security Incident Count	% of Industry Total
Acunetix	731	0.35%
Authentication bypass	5	0.00%
Burp Collaborator	89	0.04%
Directory Traversal	337	0.16%
DoS recon tool	8	0.00%
File Access	38	0.02%
File Upload	18,825	9.03%
Foundstone	350	0.17%
Heartbleed	1	0.00%
Injection	8	0.00%
Netsparker	355	0.17%
Nikto	443	0.21%
Null Byte Injection	2,005	0.96%
OpenVAS	21,166	10.15%
Pagerank	1,391	0.67%
Parosproxy	840	0.40%
Ransomware	44	0.02%
Remote Code Execution	33,286	15.96%
Shell Upload	139	0.07%
SQLi	101,706	48.77%
SSRF	72	0.03%
Unauthorized Access	528	0.25%
Web App Attack *unclassified*	9,704	4.65%
Web App Attack Recon	8,679	4.16%
White Hat	1	0.00%
XXE	6,686	3.21%
XSS	525	0.25%

APPENDIX B: CITATIONS

[1] **Verizon Digital Media Services**, "Verizon DBIR 2016: Web Application Attacks are the #1 Source of Data Breaches," **https://www.verizondigitalmedia.com/blog/2016/06/verizon-dbir-2016-web-application-attacks-are-the-1-source-of-data-breaches/**. Posted June 21, 2016.

[2] **Gartner cited in TechBeacon**, "Highlights from the 2015 Gartner Magic Quadrant for application security testing," **https://techbeacon.com/highlights-2015-gartner-magic-quadrant-application-security-testing**. Post date unknown.

[3] **SANS cited in Veracode.com**, "OWASP Top Ten Vulnerabilities," **https://www.veracode.com/directory/owasp-top-10**. Posted

[4] **Computerweekly.com**, "Veracode finds most code fails OWASP security checklist," **http://www.computerweekly.com/news/4500259915/Veracode-finds-most-web-apps-fail-Owasp-security-check-list**. Posted Dec. 13, 2015.

[5] **Washington Post**, "How a grad student trying to build the first botnet brought the Internet to its knees," **https://www.washingtonpost.com/news/the-switch/wp/2013/11/01/how-a-grad-student-trying-to-build-the-first-botnet-brought-the-internet-to-its-knees/?utm_term=.7f8a450ffc90**. Posted November 1, 2013.

[6] **Verizon 2016 Data Brach Investigations Report**, **http://www.verizonenterprise.com/resources/reports/rp_DBIR_2016_Report_en_xg.pdf**.

[7] **Mitre.org**, "CVE-2012-3414," **https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2012-3414**

[8] **Threatpost**, "New Server-Side Ransomware Hitting Hospitals," **https://threatpost.com/new-server-side-ransomware-hitting-hospitals/117059/**. Posted March 29, 2016.

[9] **US-CERT**, "Alert (TA15-119A), Top 30 Targeted High Risk Vulnerabilities," **https://www.us-cert.gov/ncas/alerts/TA15-119A**. Posted April 29, 2017 / last revised September 29, 2016.

[10] **Correspondence from Sen. Bill Nelson to CloudPets**, published on TroyHunt.com, **http://files.troyhunt.com/03.07.17%20BN%20Letter%20to%20Spiral%20Toys%20re%20Data%20Breach.pdf**. Dated March 7,2017.

[11] **Archer Security Group**, "What Happens When a Company Goes Belly-Up with Your Kids' Data?," **http://www.archersecuritygroup.com/happens-company-goes-belly-kids-data/**. Posted March 2017.

[12] **TalosIntelligence.com**, "Content-Type: Malicious: New Apache Struts2 0-Day Under Attack," **http://blog.talosintelligence.com/2017/03/apache-0-day-exploited.html**. Posted March 8, 2017.

[13] **Apache.org**, "Security bulletins: S2-045," **https://cwiki.apache.org/confluence/display/WW/S2-045**. Posted March 19,2017.

[14] **Mitre.org**, "CVE-2017-5638," **http://www.cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2017-5638**. Posted March 2017.

[15] **CBC News**, "Canada Revenue Agency shuts down online services after discovering "internet vulnerability," **http://www.cbc.ca/news/canada/cra-online-1.4021131**. Posted March 11, 2017.

[16] **CBC News**, "Federal officials sat no personal information leaked in 'credible' software security threat," **http://www.cbc.ca/news/politics/cra-internet-vulnerability-government-1.4022591**. Posted March 13, 2017.

CONTRIBUTORS

Blake Allen
Allison Armstrong
Willie Blue
Carlos Castillo
Jon Espenschied
Michael Farnum
Sean Ferguson
Misha Govshteyn
Angela Gunn
Joseph Hitchcock
Ashley Jackson
Kevin Keenan
Scott Lambert
Jonny Milliken
Audian Paxson
Ian Rickey
John Whiteside
Marc Willebeek-Lemair

Illustrator:
Louisa Clayton

ALERT LOGIC®
Security. Compliance. Cloud.

alertlogic.com

CORPORATE HEADQUARTERS

Alert Logic, Inc.
1776 Yorktown, 7th Floor
Houston, TX 77056
+1.713.484.8383

UK OFFICE

Floor 5, 1 Capital Quarter
Cardiff
CF10 4BE
United Kingdom
+44 (0) 203 011 5533